

(19)



Europäisches Patentamt

European Patent Office

Office européen des brevets



(11)

EP 0 896 279 A1

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:

10.02.1999 Bulletin 1999/06

(51) Int Cl.⁶: G06F 12/08

(21) Application number: 98305910.6

(22) Date of filing: 24.07.1998

(84) Designated Contracting States:

AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE

Designated Extension States:

AL LT LV MK RO SI

(30) Priority: 06.08.1997 US 908140

(71) Applicant: International Business Machines
Corporation

Armonk, N.Y. 10504 (US)

(72) Inventors:

- Gilda, Glenn David
Binghamton, New York 13903 (US)
- Gregor, Steven Lee
Endicott, New York 13760 (US)

(74) Representative: Jennings, Michael John

IBM United Kingdom Limited,
Intellectual Property Department,
Hursley Park
Winchester, Hampshire SO21 2JN (GB)

(54) Method and apparatus for configurable cache coherency

(57) A coherency controller for configurable caches. A base microprocessor design accommodates system configurations both with and without L2 cache tag and data arrays installed. Second level cache control logic exists within the microprocessor chip, and when the external second level cache tag and data arrays are removed their inputs to the microprocessor are tied to an inactive state. A configuration switch is set in the second level cache controller that causes snoop requests from a system bus to get reflected onto a first level cache snooping path. The first level cache status is then fed back to the second level cache controller, in a manner consistent with the timing required for support of a second level cache search, and fed into the second level cache status signal generation logic, effectively making the second level cache controller believe that the second level cache still exists for snooping. All other actions remain the same in the second level cache controller providing an effective and simple method for supporting snooping bus protocols. A result is that now every bus request snoops the first level cache without knowledge of presence of an L2 cache. This environment is provided to support entry level single processor configurations where the snooping requests only amount to input/output traffic.

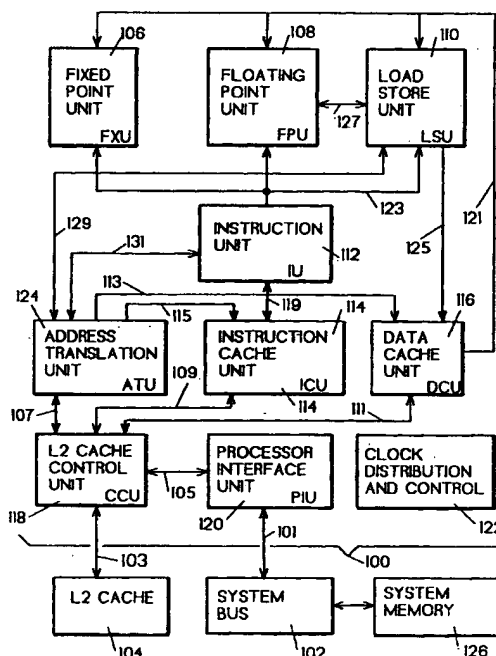


FIG. 1

BEST AVAILABLE COPY

Description

Field of the Invention

[0001] This invention generally relates to caches for computer systems, such as set associative caches and direct-mapped caches, and more particularly to providing data coherency in systems including processors having two levels of cache and those having a single level cache.

Background Art

[0002] Microprocessor development projects are costly in terms of time and money. As such, there is a strong desire to make each of these design efforts versatile enough to cover the entire planned product range. The problem stems from the fact that the entry level machines must be inexpensive in deference to performance, while the upper end of the product range must favor performance.

[0003] A common way to achieve this involves making the microprocessor core "programmable." This translates into a need to make it work with different external environments. To be specific, it is designed to accommodate multiple levels of cache buffer, with the first built into the microprocessor core, and the second external to the chip. This external second level cache can be quite expensive and unnecessary for the entry level models. However, it is absolutely necessary to the high-end which requires that multiple processors with their external second level caches are coupled together and cooperatively work on programs and data.

[0004] A cache is a high speed buffer which holds recently used memory data. Due to the locality of references nature for programs, most of the access of data may be accomplished in a cache, in which case slower accessing to bulk memory can be avoided.

[0005] A typical shared memory multiprocessor system implements a coherency mechanism for its memory subsystem. This memory subsystem contains one or more levels of cache memory associated with a local processor. These processor/cache subsystems share a bus connection to main memory. A snooping protocol is adopted where certain accesses to memory require that processor caches in the system be searched for the most recent (modified) version of requested data.

[0006] In accordance with an exemplary high-end system, a two level cache subsystem with level 2 (L2) cache line size some power of 2 larger than level 1 (L1) cache line size is implemented. Both caches implement writeback policies, and L1 is set-associative. L1 is subdivided into sublines which track which portions of the cache line contain modified data. Snoop requests from the bus are received at L2 and, if appropriate, the request is also forwarded on to L1. The snoop request forwarded to L1, however, requires accessing the L1 directory for all of the consecutive L1 cache entries which

may contain data associated with the L2 cache line.

[0007] Two fundamental ways to manage the contents of multiple levels of caches exist:

1. Allow unique: the first level and second level caches are allowed to have unique data.

2. Force inclusion: the first level of cache is required to be a subset of the second level of cache.

[0008] In a coherent shared memory multiprocessing environment, each time a processor issues a request for memory data, the other processors' caches may need to be searched for copies of that data, depending upon the type of request. Also, in a system with a single processor, memory coherency needs to be maintained with other devices that access memory, such as I/O processors.

[0009] Consider a first exemplary system including a plurality of high-level processors -- such that private first and second level caches exist for each processor in the multiprocessing system. Assume, also that a snooping bus protocol is used for the multiprocessor memory hierarchy and that cache blocks may exist in cache in the modified, exclusive, shared, or invalid state (MESI protocol).

1. ALLOW UNIQUE: In a first exemplary system operating in an "allow unique" environment where the first level and second level caches are allowed to have unique data, inasmuch as data may exist in either one or both levels of cache within the individual processors, maintaining cache coherency requires that both the first and second level caches be searched each time an alternate processor request shows up on the snooping bus. This causes interference at the first level cache which may adversely affect performance of the executing instruction stream within the processor. Conversely, it requires an added port on the first level cache to allow interrogation of the cache for snooping requests. In any case, one is faced with pulling data from one or both levels of cache in order to supply it to requests from other processors. Additionally, coordination of the snoop responses between the two levels of cache is required within each processor.

Consider a second exemplary system including one or more entry-level processors -- such that private first level caches exist for all processors in the multiprocessing system, and second level caches do not exist for all processors in the system. In an "allow unique" environment the first level cache already has the mechanisms in place to allow snoop requests to be handled.

2. FORCE INCLUSION: In a forced inclusion environment, the contents of the first level cache must always be a subset of the second level cache within

an individual processor. This allows a system to be built where the second level cache can control the snooping of external bus requests without interference in the first level cache unless data actually exist there that are required. This is facilitated by a first level cache status array maintained at the second level cache controller which indicates presence, and possibly state, of the first level cache data. Note that presence is minimally required as the state of the first level data must be less than or equal to that of the encompassing second level cache block. Also, in this exemplary system, the L2 cache controller has primary responsibility for managing cache coherency. It also forms requests for data and other commands based on processor requests which are forwarded to the system bus.

[0010] In a forced inclusion environment, for entry-level microprocessor where the second level cache is not included, the ability to control cache coherency is more difficult because the mechanisms for this control reside in the L2 cache controller.

[0011] There is a requirement for a cooperative work environment in multiprocessor configurations including either low-end or high-end microprocessors. There is also a requirement to maintain data coherency in such a system.

[0012] There is also a requirement to provide a microprocessor design which is adaptable to either the low-end or high-end configuration.

Summary of the Invention

[0013] In accordance with a first aspect of the invention, an apparatus and method is provided for operating a computing system including an L1 cache. A second level cache controller includes a first level cache snooping path and an interface to a system bus. The controller is operable in the absence of a second level cache for reflecting snoop requests received at the system bus interface and requests received at a processor interface for a second level cache to the first level cache snooping path.

[0014] In a second aspect of the invention, a cache controller for the management of contents of a single level of cache includes an interface to a processor, and a second level cache controller operable in the absence of a second level cache for reflecting processor snoop requests received at the processor interface to a first level cache snooping path.

[0015] Other features and advantages of this invention will become apparent from the following detailed description of the presently preferred embodiment of the invention, taken in conjunction with the accompanying drawings.

Brief Description of the Drawings

[0016] Figure 1 is a block diagram illustrating a typical microprocessor architecture within which a preferred embodiment of the invention is implemented.

[0017] Figure 2 illustrates how Figure 2A through 2C relate, while the latter are block diagrams showing the implementation of a preferred embodiment of the invention within the microprocessor of Figure 1.

[0018] Figures 3-6 are block diagrams illustrating the system and L2 cache bus interfaces 101 and 103 of Figure 1, with Figure 3 generally illustrating the system data bus; Figure 4, the system bus controls; Figure 5, the L2 cache data bus; and Figure 6, the L2 cache controls.

[0019] Figure 7 is a block diagram illustrating those portions of the L2 cache data bus controls of Figure 5 utilized in the entry-level microprocessor design of the preferred embodiment of the invention. The unused portions are shown hatched out.

[0020] Figure 8 is a block diagram illustrating those portions of the L2 cache controls of Figure 6 utilized in the entry-level microprocessor design of the preferred embodiment of the invention. The unused portions are shown hatched out.

[0021] Figure 9 is a state diagram illustrating the tag state machine operated by the L2 cache controller of Figure 2C when responding to a request.

[0022] Figure 10 is a flow diagram illustrating the four stages implementing a pipelined L1 snoop operation, and also the search operation of the invention.

[0023] Figure 11A and 11B are system diagrams illustrating systems including processors with and without an L2 cache in accordance with the invention.

[0024] Figure 12 is a timing diagram illustrating the timing of control and status signals on the system bus for explicit search (snoop) operations.

[0025] Figure 13 is a timing diagram illustrating the timing of control and status signals on the system bus for implied search operations.

Detailed Description of Preferred Embodiments

[0026] As used in the art, the term "snoop" may have several meanings. For the purpose of following description of the invention, the term "bus snoop" refers to L1 and/or L2 cache directory access operations originating from a system bus request: this will be at the L2 cache and possibly at the L1 cache in L2 cache installed mode, or at L1 cache in L2 not installed mode. The term "L1 snoop" refers to an access operation occurring at the L1 cache directory from either a system bus or processor request. A "bus snoop" can result in an "L1 snoop". The term "processor snoop" refers to a snoop of L1 cache generated by a processor request which originates on the interface to the L2 cache controller from either the ATU 124 or DC 116; this request also accesses the L2 cache directory in systems with L2 installed. A "processor snoop" can result in an "L1 snoop".

Part 1

[0027] Referring to Figure 1, the microprocessor architecture within which a preferred embodiment of the invention is implemented will be described.

[0028] Microprocessor chip 100 is organized to interface system bus 102 and L2 cache 104, and includes the following functional units: fixed point unit (FXU) 106, floating point unit (FPU) 108, load store unit (LSU) 110, instruction unit (IU) 112, instruction cache unit (ICU) 114, data cache unit (DCU) 116, L2 cache control unit 118, processor interface unit (PIU) 120, clock distribution and control 122, and address translation unit (ATU) 124. In a multiprocessor environment, several processors 100 and their associated L2 caches 104 may interface system bus 102 over buses equivalent to bus 101, and share access through system bus 102 bus to main memory (sometimes referred to as L3 memory) 126.

[0029] The various functional units of microprocessor 100 interface over data, address, and/or control I/O pins, lines and/or busses as will be described hereafter. When referring to a figure, "line" can refer to either a single signal line or a collection of signal lines (i.e., a bus). Those functional units most pertinent to the invention, and which will be described in greater detail hereafter, include the load/store unit (LSU) 110, the data cache unit (DCU) 116, the L2 cache control unit (CCU) 118, and the address translation unit (ATU) 124.

[0030] In broad overview, the functional units on chip 100 communicate as follows. Clock distribution and control 122 provides clocking signals to all functional units on microprocessor chip 100. System bus 102 interfaces to PIU 120 over bidirectional bus 101, and thence over buses 105 with CCU 118. L2 cache 104 communicates with CCU 118 over buses 103. CCU 118 communicates instructions with ICU 114 over buses 109, with DCU 116 over buses 111, and provides address information to ATU 124 and receives miss interface signals over buses 107. LSU 110 and IU 112 provide request interfaces to ATU 124 and receive translation state information over lines 129 and 131. ATU 124 provides translated address to ICU 114 over lines 115, and to DCU 116 over lines 113. ICU 114 interfaces to instruction unit 112 over bus 119. DCU 116 provides data to FXU 106, FPU 108 and LSU 110 over bus 121, and IU 112 provides instructions to FXU 106, FPU 108 and LSU 110 over bus 123. LSU 110 provides data to DCU 116 over bus 125. FPU 108 provides and receives data to DCU 116 over buses 127 to LSU 110, then across buses 125. Processor 100 accesses main memory through system bus 102.

Microprocessor Core 100

[0031] Referring to Figures 2A through 2C, and Figures 3-6, the core of microprocessor 100 will be described. Figure 2A generally corresponds to load/store unit (LSU) 110, Figure 2B to address translation unit (ATU) 124, and Figure 2C to data cache unit (DCU) 116.

Figures 3-6 generally correspond to L2 cache control unit (CCU) 118 and processor interface unit (PIU) 120.

[0032] Dispatch block 300 directs instructions from instruction unit 112 to the DECODE stage buffers of the various execution units 106, 108, 110, including on bus 301 (which is that portion of buses 123 directed to LSU 110) to LSU pipeline buffer 302.

[0033] The function of load/store unit 110 is to generate effective addresses on 64 bit wide bus 313 for load and store instructions and to serve as a source and sink for GPR data. During writes to cache 400, registers 314 and 316 hold the data and address, respectively; the effective address is on bus 313, and data select block 320 puts the data out on bus 323. During cache reads, data from cache 400 comes in on line 461, is latched in register 330, and from there sent on line 333 to general purpose registers 306 or to fixed point unit 106.

[0034] The output of LSU pipeline buffer 302 is fed on line 303 to the LSU decode and address generation block AGEN 304, which contains the general purpose registers 306 and address generation adders (not shown). The data output of decode block 304 is fed on lines 311 to data register 314 and thence on line 319 to data select block 320. The address output of AGEN 304 is fed on lines 313 to EXECUTE stage buffer 316, and on bus 309 to real address MRU 430. AGEN 304 output also includes control line 307, which it sets to indicate either real or virtual mode addressing to data cache control block 470 (also referred to as data cache controller or first level cache controller.)

[0035] The outputs of buffer 316 are fed on lines 317 to data select block 320 and to data cache address register 408, DIR address register 414 and register slot MRU address register 406. The output of register 408 is fed on line 409 to multiplexer 412. Data select block 320 contains the data to be stored to data cache 400 from load store unit 110, and this is fed thereto on store data output lines 323 via multiplexer 432, lines 433, align block 460, lines 461, register 456, lines 457, and line 427 via multiplexer 426. Data select block 320 also provides control signals to data cache controller 470 on lines 321. The other inputs to multiplexer 432 are (1) L2 corrected data 609 via multiplexer 426 and line 427, which is also fed to data cache 400, (2) bypass data to DC on line 621, and (3) unaligned data (aka store merging and correction) register 452 via lines 453 to line 427 via multiplexer 426. Multiplexer 432 output line 433 is also fed via align block 460 and line 461 to register 456 and thence via multiplexer 424 to L2 cache controller on line 425, along with the output of castout buffer 450 on line 451. Align block 460 is, in this embodiment, a barrel rotator or shifter which aligns D cache 400 data to quad word boundaries on reads, and from multiplexer 432 to quad word boundaries on stores.

[0036] An effective address from instruction unit 112 on line 367 (a portion lines 131) is latched in register 364 and fed on line 365 to ITLB 358 and to the compare and address select block 356 at ISLB 354. Line 313 from

AGEN 304 is latched in register 384, and fed on line 385 to DTLB array 378 and compare and address select block 374 at DSLB 376. In this preferred embodiment, DTLB 378 may be a standard design, such as that described by Liu, supra. Whereas the Liu TLB design is 32 bits wide, in this preferred embodiment a 64 bit wide TLB 378 is used.

[0037] Data select 320 output on line 325 is fed to PUTAWAY stage buffer 330, which also receives data on lines 461 from data cache 400 (via lines 401 and align block 460) for LSU 110, and FPU 108 results on line 327 which is a portion of bus 127. The output of PUTAWAY stage buffer 330 is fed on lines 333 to a floating point register in FPU 108, special purpose registers 334 (among which are the timers), and general purpose registers 306. Special purpose registers 334 output line 335 is fed back to data select block 320 which allows the processor to read them. Line 333 carries the data for FPU 108 when doing a fetch from cache 400.

[0038] The selected output of instruction segment look aside buffer (ISLB) 354 is fed on lines 355 to comparator 362, along with the virtual address output of ITLB 358 on lines 359. ITLB real address output lines 361 are input to IC controls 350 (which includes instruction directory 352) and status information on line 361 is fed to ATU controls 370. The output of comparator 362 is fed on lines 363 to IC controls 350 and to ATU controls 370. The output of DSLB 376 is fed on lines 377 to comparator 382, along with the output of DTLB 378 on lines 379. The output of comparator 382 is fed on lines 383 to ATU controls 370 and DC control 470. DTLB 378 status output 381 is fed to ATU controls 370 and DC control 470. ATU controls 370 outputs include control lines 369 to data cache controller 470, L2 address 371 and ATU write data 373. IC controls 350 output is L2 address line 351. Real address from DTLB 378 is fed on lines 381 to DC address register 408 and directory address register 414.

[0039] Predicted real address MRU 430 output signals on line 431, representing the predicted read address bits 50:51, are latched in registers 410 and 416. The output of data cache address register 410 on line 411 is multiplexed with bits 50:51 of the output of register 408 in multiplexer 412, and its output is fed on address lines 413 to data cache 400. The remaining bits of DC address register 408 are fed straight through on line 413 to data cache 400. Similarly, the output of register 416 is fed on lines 417 to multiplexer 436, where it is multiplexed with bits 50:51 of the output of register 414 on line 415, and the result fed on lines 437 to directory array 440. The output of register 414 on line 415 is also fed to address register 408.

[0040] The function of real address MRU 430 is to provide predicted real address bits 50:51 to data cache 400 and directory array 440.

[0041] During the fetch stage, data cache 400 output 401 is fed to unaligned data register 452 and align block 460, and thence on line 461 to registers 456 and 330.

Line 401 contains the data to be read from data cache 400 by the load store unit 110, L1 snoop data to the L2 cache controller 118, merge data for partial stores to the data cache 400, and castout data to castout buffer 450. Slot MRU 402 output line 403 controls the selection of one of four sets of data to load on bus 401 through a multiplexer (not shown) on the output of data cache 400.

[0042] The output of castout buffer 450 is multiplexed in multiplexer 424 with the output of register 452 on lines 453 and line 457 from DC putaway register 456, the output appearing on lines 425 to the L2 cache controller. The output of register 452 along with DC putaway register 456 and L2 corrected data on line 609 is also fed to data cache input multiplexer 426, the output of which appears on lines 427 to data cache 400 and multiplexer 432. The output of register 406 is fed on line 407 to slot MRU 402. Slot MRU 402 output 403 is fed to data cache 400 where it controls a data multiplexer which selects the appropriate cache set (as taught by Liu, supra.)

[0043] Data cache (DC) control 470 receives inputs from directory array 440 on lines 441 (signifying a directory array hit or miss), from AGEN 304 on lines 307, data select and execute cycle control block 320 on lines 321, ATU controls 370 on lines 369, comparator 382 on lines 383, and controls 660 on search control line 481. Its outputs are L2 address line 471, true slot hit or castout slot line 849 (which is fed to control block 660 and is used for maintaining the L1 status array in the L2 cache controller), cache miss line 294 to processor 100, and status response line 483, which also includes a cache hit/miss indication for search operations, to L2 cache control block 660.

[0044] The function of data cache control 470 is to control the data flow multiplexing into and out of data cache 400 and send results to the load/store unit 110, address translation unit 124, and L2 cache control unit 118, and also to control writing of data into data cache 400.

[0045] Data directory 440 contains address tags to indicate if the contents of the real address are present in cache 400, and the status of the cache lines, whether modified, shared, or invalid. It also contains an LRU pointer for each congruence class, indicating which cache 400 line should be replaced.

[0046] Address translation unit (ATU) control 370 handles translations from effective addresses to virtual addresses to real addresses. It receives as inputs L2 corrected data on line 353, and provides TLB reload data output on lines 375 to instruction translation lookaside buffer (ITLB) 358 and data translation lookaside buffer (DTLB) 378, ISLB 354, and DSLB 376. With respect to look aside tables 354, 358, 376, 378, if a miss condition is detected, ATU sequencer 370 requests data (address and length) to L2 cache on bus 371 (Fig. 6.) When L2 responds on bus 353 (Fig. 5), ATU examines the data to select data for look aside buffer 378, 376, 354, 358, as the case may be, or signals a translation exception back to the instruction unit. ATU controls 370 tracks seg-

ment and page table updates and sends them to L2 controls on line 371. Line 381 provides the real address to the data cache directory for comparison.

[0047] The effective address is compared in ISLB 354 comparator 356 with the virtual address. If these match, then a valid effective to virtual address translation exists in buffer 354, which transmits the virtual address on line 355 to compare block 362.

[0048] ITLB 358 is accessed by an effective address on line 365 from register 364 for doing virtual to real address translation. The address input to ITLB 358 is a portion of the effective address from IU 112 on lines 367. Comparator 362 compares virtual addresses on lines 355 and 359, and signals the result on line 363. Associated with each virtual address in ITLB array 358 is a real address. The signal on line 363 indicates whether or not the address on line 361 is valid.

[0049] DTLB 378 is accessed by an address from register 384. Comparator 382 compares data on lines 379 and 377, and signals the result on line 383. The signal on line 383 indicates whether or not the address on line 379 is valid.

System Bus Interface 120

[0050] Referring to Figures 3 through 6, the system bus interface 120 and L2 cache control unit 118 of Figure 1 will be further described.

[0051] Correspondence between the high level block diagram of Figure 1, and the more detailed illustration of the preferred embodiment in Figures 3 to 6, is as follows. Bus 101 of Figure 1 corresponds in Figures 3-6 to system controls lines 559 at driver/receiver 556, system address lines 569 at driver/receiver 564, system data hi bus 513 at driver/receiver 512, and system data low bus 517 at driver/receiver 516. Bus 103 to L2 cache 104 of Figure 1 corresponds to L2 cache address lines 691 out of driver 690, L2 tag address line 693 out of driver 692, L2 tag data lines 697 at driver/receiver 694, and L2 cache data bus 645 at driver/receiver 644. ICU bus 109 of Figure 1 corresponds (from ICU) to IC request lines 351, and (to ICU) DOIC register 606 output lines 607, and bypass to IC multiplexer 616 on lines 617. DCU bus 111 of Figure 1 corresponds (from DCU) to DC request lines 471 and data cache write data bus 425, and (to DCU) to bypass to DC multiplexer 620 on lines 621 and data cache data out (DODC) register 608 output line 609. Address translation unit (ATU) input/output bus 107 of Figure 1 corresponds to ATU request lines 371, ATU write data bus 373, and multiplexer 612 output lines 353.

[0052] Referring to Figures 4 and 6, requests to L2 cache control 118 are latched in address/command register 650 from ATU request lines 371, IC request lines 351, DC request lines 471, and on lines 567 from address in register 566, which latches system bus addresses on lines 565 from receiver 564. These address/command signals are latched as required in registers 650, 652 and 654 connected via lines 651 and 653. The

output of the third register 654 is fed to controls block 660 on line 655. The output of first stage register 650 is fed on lines 651 to register 652, driver 690 to provide L2 cache address signal 691, driver 692 to provide L2 tag address signal 695, ECC checking circuit 684, address comparator 664, controls block 660, cache controller (CC) snoop address register 670, processor address registers CBPADR 674 and CBMADR 676, and address multiplexer 680. ECC 684 output is fed on lines 685 to driver 694 to provide L2 tag data on lines 697. CBPADR address register 674 contains the address to the system bus in the event of a cache miss, the output of which is fed to multiplexer 680 on line 675. CBMADR address register 676 contains the snoop address portion, and its output is fed to multiplexer 680 on line 677. Receiver 694 output from L2 tag data lines 697 is fed on lines 695 to L2 tag in register (L2TAGIN) 688 and thence on lines 689 to error correction code (ECC) block 686. The output of ECC block 686 is fed on lines 687 to comparator 664, address registers 670, 674 and 676.

[0053] The output of comparator 664 is fed on line 665 to controls block 660. CCS address register 670 output line 671 is fed to multiplexer 678 along with lines 651 to generate the data cache L1 snoop address on lines 679.

[0054] The output of address out multiplexer 680 is fed on lines 681 to address out register 560, and thence on line 561 to the system address bus 569 through driver 564. The output of controls block 660 is fed on lines 663 to arbitration and control block 552, and on lines 661 to address/command register 658. Arbitration and control block 552 receives control data from receiver 556 via lines 557, and provides output on lines 555 to controls block 660, and in the event of an L2 cache miss request out control signals are sent on line 553 through driver 556 to system controls bus 559. Another output of controls block appears on lines 661 to address/command register 658, the output of which appears on line 659 to multiplexer 672. Multiplexer 672 also receives input from lines 653 and 655, and provides its output on lines 673 back to register 650.

[0055] Referring to Figure 5, ECC block 632, DOIC register 606, DODC register 608, L2PDO register 636, multiplexer 616 and multiplexer 620 each receive inputs from data input register 624 on bus 625. The output of ECC block 632 is fed on line 633 to L2 data out register 638, and thence to driver 644 on line 639. The output of L2PDO register 636 is fed on line 637 to inpage buffer 646, the output of which is fed on line 647 to L2PDI register 642 and ECC circuit 632. The output of L2PDI register 642 is fed on line 643 to DOIC register 606, DODC register 608, CCDI register 624, and to bypass multiplexers 620 and 616. The output of multiplexers 620 and 616 represent bypass data, and are fed on lines 621 and 617 to the DC and IC, respectively. Data cache write data line 425 is fed to CMCD register 628 and CCDI register 624. The output of CMCD register 628 is fed on lines 629 to L2PDO register 636, and castout buffers 602.

[0056] Referring to Figures 3 and 5, L2 cache data in from bus 645 is received at receivers 644, fed on line 649 to L2 data in register 640 and thence on lines 641 to ECC circuitry 634 and bypass multiplexers 616 and 620. From ECC circuitry 634, L2 cache data in is fed on lines 635 to cache controller data in register (CCDI) 624, DOIC register 606 and DODC register 608. DODC register 608 output 609 is fed to data cache unit 116 (Figure 1), DC bypass multiplexer 620, ATU multiplexer 612, and castout buffers 602. The output of DOIC register 606 is fed on lines 607 to instruction cache unit 114 (Figure 1), ATU multiplexer 612, and castout buffers 602. Castout buffers 602 output on lines 603 is fed to data high output register 502 and multiplexer 520, the output of which is fed on lines 521 to data output registers 502 and 504.

[0057] In operation, registers 624 and 636 form a pipeline buffer to inpage buffer 646 and register 642. Inpage buffer 646 caches a line from the system bus. Line 641 from L2 data in register 640 to bypass multiplexers 616, 620 allows the saving of a cycle on cache misses when error correction is not required. DOIC register 606 provides corrected data to instruction cache unit 114, and DODC provides corrected data to data cache unit 116. Either register may supply data to the ATU 124.

[0058] The normal path for routing L2 cache data is through register 640, ECC 634, and DOIC register 606 and DODC register 608.

Processor Interface Unit 120

[0059] Referring now to Figure 3, a more detailed description of processor interface unit 120 of Figure 1, and associated circuitry, will be provided. Figure 3 represents the data flow portion of PIU 120 and System Bus 102.

[0060] System bus 102 data high bus 513 and data low bus 517 communicate through driver/receivers 512 and 516, respectively with data high output register 502 on lines 503, data high in register 506 on lines 515, data low out register 504 on lines 505, and data low input register 508 on lines 519. Each of busses 513, 517 is capable of handling eight bytes of data, providing a 16 byte data bus. If the system is operating on only eight bytes, only one set of the input/output registers (such as 504, 508) is used.

[0061] System data input registers 508 outputs on lines 507 and 509, respectively, are fed to multiplexer 524 and thence, along with registers 506 on lines 507, on lines 525 to cache control data in (CCDI) register 624 (Figure 5), which is the main data input register of the cache controller. Data input register 624 output is fed on bus 625 to multiplexer 520.

Load/Store Unit (LSU) 110

[0062] Load/store unit (LSU) 110 functions to decode fixed point and floating point loads and store and cache

management operations, and to send effective addresses and storage commands to the data cache unit (DCU) 116. LSU 110 also handles most move-to and move-from special purpose register (SPR) 334 instructions. In addition to functioning as a load/store unit, LSU 110 also controls instruction execution sequencing after instructions have been dispatched, through detection of most instruction execution interlocks, and the generation of resulting pipeline hold signals.

[0063] LSU 110 provides a six port register file 306, made up of four 32x18 register array macros, arranged as a 32x72 array with two write ports and four read ports. This array implements the 64-bit general purpose registers (GPRs) 306. GPR array 306 also provides operands for fixed point unit (FXU) 106 decode stage (not shown) as well as for LSU 110. FXU 106 decodes its own instructions and generates requests to LSU 110 for the necessary operands, as well as providing on line 327 a result operand and address when appropriate. LSU 110 accesses GPRs 306 for registers needed to generate effective addresses (EA), and for data for store instructions. Data operands received from data cache 116 on line 461, and updated effective addresses are written back to the GPRs by LSU 110. Lines 327 contain the FPU results, and are fed to register 330.

[0064] In handling floating point loads and stores, LSU 110 generates effective addresses using operands from GPR 306, and accesses the necessary floating point register (FPR) operands from the floating point unit (FPU) 108.

[0065] Instructions dispatched to LSU 110 are latched in its DECODE cycle instruction register 302 at the end of the I- fetch cycle. The basic LSU 110 pipe is three stages: DECODE 302/304, EXECUTE 316/320, and PUTAWAY 330. During the DECODE cycle corresponding to 302/304, the instructions are decoded, and operands are fetched from the GPR 306 array. Addressing operands are gated to a 64-bit address generation (AGEN) adder, and a 64-bit effective address is calculated. The effective address (EA) is sent on lines 313 to the address translation unit (ATU) 124 and to data cache unit (DCU) 116 and latched at the end of the DECODE cycle in pipeline buffer 316 which holds the effective address during the EXECUTE cycle.

[0066] During the EXECUTE cycle, the operand for store operations is passed to the DCU on line 323, where it is aligned in block 460 and saved in register 456 for PUTAWAY in D-cache 400. At the end of the EXECUTE cycle, if a load type instruction is being executed, the data operand returns on line 461 to LSU 110 from the DCU, and is saved in pipeline buffer 330 for PUTAWAY.

[0067] During PUTAWAY cycle 330, as is represented by lines 333, up to two 8-byte or one 16-byte operand may be written to GPR 306. Floating point loads are limited to one 8-byte operand per cycle. GPRs 306 are not written until late in the PUTAWAY cycle 330. This requires that operands being written to these arrays be

bypassed around the arrays to reduce pipeline interlocks. Delaying the write to GPR 306 also allows sign extension for algebraic load operations to be performed, helping to balance path delays between EXECUTE 316/320 and PUTAWAY 330 cycles for these instructions.

Fixed Point Unit (FXU) 106

[0068] Fixed Point Unit (FXU) 106 executes the fixed point instructions, not including storage access instructions. FXU 106 includes a 64-bit adder, a 64-bit logical unit, a 64-bit rotate-merge unit, and a 64-bit carry save adder which supports two-bit-per-cycle product formation during multiply instructions.

[0069] During division, quotient formation occurs one bit per cycle, through repeated subtraction of the divisor from the shifted dividend.

Floating Point Unit (FPU) 108

[0070] Floating point unit (FPU) 108 executes the floating point instructions, but not the storage access instructions. In one exemplary embodiment, FPU 108 includes a 5-port 32x72-bit register array, a 32-bit status-control register, a 3-bit overlap scan booth encoder unit, 2-bit quotient generation unit, a 106-bit carry save adder, a 106-bit increment-full adder, an operand alignment shifter unit, a normalizer unit, and a rounder unit.

Address Translation Unit (ATU) 124

[0071] Referring, primarily, to Figure 2B, address translation unit (ATU) 124 translates the data effective address (EA) from load/store unit (LSU) 110 and the instruction effective address from instruction unit 112 into real addresses used by the Data and Instruction Caches to access their L1 Caches and used by the L2 Cache Control Unit 118 to access the L2 Cache 104.

[0072] Microprocessor 100 implements segment lookaside buffers (SLB) 354, 376 and translation lookaside buffers (TLB) 358, 378, which function as caches for segment and page table entries. When a required entry is not found in a look aside buffer, ATU 124 initiates a fetch to L2 cache control 118 to access segment and page table entries from memory 126 or L2 cache 104.

[0073] ATU 124 reports any translation data storage interrupts (DSI) to the load/store unit 110 and any translation instruction interrupts to the instruction unit 112. Reference, change and tag change bits are all updated by store requests to cache control 118 from ATU 124.

[0074] Microprocessor 100 provides a 4-entry SLB 354 for instruction address translation and an 8-entry SLB 376 for data address translation. SLBs 354, 376 contain the most-recently translated segments in a fully associative arrangement. The ESID (Effective Segment ID) portion of the effective data or instruction address is compared 356, 374 simultaneously to all entries in the

respective SLB 354, 376 ESIDs when segment translation is enabled.

[0075] ATU 124 includes separate instruction and data TLBs 358, 378, respectively, to hold the results of virtual to real address translations. With virtual to real translation active, the VSID from the matching SLB 354, 376 is compared in comparators 362, 382 against the VSID stored in the TLB 358, 378. If a compare is found, the Real Page Number (RPN) stored in the matching TLB 358, 378 entry is used to form the real address. Replacement is managed independently in each TLB 358, 378 by an LRU bit for each of the 256 pairs of entries.

L1 Data Cache Unit (DCU) 116

[0076] In a preferred embodiment, L1 data cache unit (DCU) 116 has the following attributes: 64 KB size, 64 byte line size, 4-way set-associative, 2 subline-modified bits per line, MRU slot selection, 40-bit real address, 16-byte dataflow to/from processor, store-in design, and multi-processor support. The term "cache line" refers to a 64-byte block of data in the cache which corresponds to a single cache directory entry. Slot MRU 402 provides selection of one of four sets of cache data during an execute cycle. Real address MRU 430 supplies bits 50:51 to cache 400 and cache directory 440. Error correction ECC (not shown) is provided on cache 400 and cache directory 440. Write-thru mode is implemented.

[0077] The data cache 116 array 400, representing a collection of sub arrays, is based on a 1024x78 1R1W "virtual" 2-port array macro. It provides for a read operation followed by a write operation within a processor cycle. Read data remains valid on the array outputs until the next read operation begins even if there is an intervening write operation. Eight of these arrays are used to form a 64KB cache 400. Two arrays are used per slot to form a 16-byte dataflow in and out of the array, represented by lines 401. Data parity is stored in the array. The last bit stored in the array is odd address parity across bits 50:59 of the address used to access the data cache.

[0078] Two arrays are needed to implement data cache directory 440. The directory implements a 28-bit real page number (RPN) along with five ECC check bits. A valid and two subline modified status bits are maintained, and three check bits are stored with them. The RPN and status fields are replicated four times to represent the four sets accessed at a particular directory array address. A 3-bit LRU is shared between two directory arrays to indicate the least recently used slot.

[0079] Slot MRU 402 logically appears as a 1024x4 array where each entry is associated with a cache line in data cache 400. Bits 48:51 of the 48:57 used to access the logical array 400 are effective address bits. MRU 402 bits are updated whenever an incorrect slot guess or a cache miss occurs.

[0080] Real address (RA) MRU 430 is used to gener-

ate a prediction of real address bits 50 and 51 for addressing both cache 400 and cache directories 440. As is represented by line 309, array 430 is read as part of the AGEN stage 304 of the pipeline. If a load/store unit pipeline EXECUTE stage or latched PUTAWAY stage hold is present, then the array 430 output is not used. Real mode is also used to determine if the array 430 (Figure 2C) output is used. Real mode determination occurs in AGEN 304 which sets control line 307 to either the real or virtual mode addressing. If real mode is determined, then load/store effective address (LSEA) 317 bits 50:51 are used by register 408 to access cache 400 and register 414 to access cache directory 440 instead of RA MRU array 430 output.

[0081] Real address (RA) MRU array 430 is updated from DC address register 408 via lines 281 whenever a wrong prediction occurs with respect to translated address bits 50:51. Also, data cache address register 408 and data cache directory address register 414 are updated with proper values of address bits 50:51 via line 381 for reaccessing the cache 400 and cache directory 440 arrays. Multiplexer 412 is then switched under control of data cache control block 470 so that address register 408 is used to access cache array 400. A similar function occurs with multiplexer 436 so that register 414 is used to access the directory array 440. The LSU 110 pipeline is stalled for one cycle to allow cache 400 and directory 440 to be reaccessed in parallel in the same cycle. Data is then returned to LSU 110 via line 461 in the following cycle.

Instruction Cache Unit (ICU) 114

[0082] Instruction Cache Unit (ICU) 114 contains the physical arrays, address compares, and error checking circuitry to provide a 64KB 4-way associative instruction cache with single-bit error detection and recovery. The single-cycle cache access provides up to four instructions from a selected 128-byte cache line. Instruction cache unit 114 provides instructions to other functional units, including branch prediction.

L2 Cache Control Unit 118

[0083] The functions of the L2 cache control unit 118 are to provide processor 100 with access to a private L2 cache 104, plus access to memory 126 through system bus 102 which also supports memory coherence control for multiprocessor operations. L2 cache 104 is implemented as external static RAMs, with one set of SRAMs for the directory and another set for the data.

[0084] CCU 118 accepts commands from four sources: data cache unit 116, instruction cache unit 114, address translation unit 124, and system bus 102 via Processor Interface Unit (PIU) 120. To handle these commands, CCU 118 uses the buffer structure shown in Figure 6. External and internal commands are prioritized by CCU controls 660 and placed into ADR/CMD buffer

650. ADR/CMD buffer 650 output 651 is then used to access an L2 directory (not shown) via interface lines 693 driven by driver circuits 692 to determine the hit/miss status. Additionally, appropriate address bits from bus 651 are concurrently used to access an L1 status array (not shown) in controls 660 to determine if a data cache L1 snoop needs to be done. Finally, ADR/CMD buffer 650 is used to control updating status and tag information in the L2 directory as required, a process well understood in the art. The four L2 hit/miss states are:

1) Modified

This line is different from memory and no other coherent cache has a copy of this line.

2) Exclusive

This line is the same as memory and no other coherent cache has a copy of this line.

3) Shared

This line is the same as memory and other caches may have a copy of this line.

4) Invalid

This cache and this processor's data cache do not have a copy of this line.

[0085] Data can be in the data cache only if it is also in the L2 cache.

[0086] Commands only stay in ADR/CMD buffer 650 for three cycles, at which time the command moves to ADR/CMD buffer 652 or ADR/CMD buffer 658. A processor command is moved into the ADR/CMD buffer 652 when said command is in ADR/CMD buffer 650 and the resources it needs, such as the data flow, are not available. The command will stay in ADR/CMD buffer 652 until the resource becomes available.

[0087] Commands are moved to the ADR/CMD buffer 658 from ADR/CMD buffer 650 by way of controls block 660 when a system bus snoop command needs to use the data path. The command will stay in ADR/CMD buffer 658 until the data path is available. Commands that need to issue address commands on the system bus are placed in ADR/CMD buffer 654. The command will stay in ADR/CMD buffer 654, being retried if necessary, until a successful address status and response is received from system bus 102. If data movement is required the command is then turned over to the CCU data flow logic.

[0088] Feedback from ADR/CMD buffer 658 to ADR/CMD buffer 650 is necessary for two separate functional operations. The first feedback case is for processor read commands that encountered a shared address response from system bus 102. When the processor read command is first in the ADR/CMD buffer 650 the L2 directory is marked Exclusive, assuming that this L2 will have the only copy of the data. If another device indicates that it also has a copy of this data, by a shared

address response on system bus 102, then the L2 directory must be changed from Exclusive to Shared.

[0089] The second feedback operation is used for processor write operations that must wait for a successful system bus 102 address status and response before the data can be written. For processor stores or data-cache-block-zero (dcbz) instructions that hit shared in the L2 directory, the processor must make sure that it holds the line in the exclusive state before it updates the data. Before the processor can get ownership of the shared line it may lose the line to another device, so the feedback path is provided to reinitiate the directory access.

L2 Cache Control Unit Snoop Operations

[0090] Bus snoop commands from system bus 102 come in through processor interface unit 120 and are presented to ADR/CMD buffer 650 via bus 567. At the same time a shift register (not shown) is started. The shift register is used to 'time out' the bus snoop command. Bus snoop commands require a response within a fixed time, but the command may be delayed before being brought into ADR/CMD buffer 650 because of other higher priority commands. If the shift register 'times out', an address retry response will be issued to the system bus 102.

[0091] When a bus snoop command is accepted into ADR/CMD buffer 650 the L2 directory and L1 status array are checked. If the command hits in the L2 directory and the L1 status array, then a L1 snoop command is issued to the data cache. If data must be moved to complete the bus snoop command, it will be first moved out of the L2 cache into the castout buffer 602. Then if the data cache has a modified copy of the data, its copy of the data will be moved to the castout buffer 602 and subsequently via bus 603 to system bus 102.

[0092] The memory management policy is such that segment and page translation table entries may not be accessed directly from the L1 data cache by the ATU 124. Consequently, another type of snoop operation, a processor snoop, is done for ATU commands. When an ATU command comes in, the data cache is snooped using the L1 status array. If the data cache has modified data, the ATU command is stopped until the data is moved from the data cache to the L2 data RAMs.

Processor Interface Unit (PIU)/Bus Interface Unit (BIU) 120

[0093] Referring to Figures 1 and 3, Processor interface unit (PIU) 120 controls and monitors all communications with the main system bus 102. The main functions of PIU 120 are:

- 1) Transport commands, address, and data between CCU 118 and system bus 102.

- 2) Prune out incoming command-address transfers that do not require the attention of CCU 118.

- 3) Compensate for clock domain differences between the processor 100 units and 6xx Bus 102.

- 4) Maintain and monitor system checkstop logic for Processor Run-Time Diagnostics (PRD).

[0094] System bus interface, or processor interface unit (PIU) 120, in general, receives commands from L2 cache controller (CCU) 118 on lines 663, transforms them in block 552 to the system bus clock domain and presents them on lines 559 to bus 102. It then monitors status and response information received on lines 559 for the command and informs CCU 118 on lines 555. As commands arrive from the bus on lines 559, PIU 120 categorizes them into one of three categories: master operations, bus snoop operations and other operations. Master operations are those originated by CCU 118 on the same chip 100 as PIU 120. These operations need to be monitored for status and response, updating CCU 118 as this information arrives. Bus snoop operations are those that are originated by other bus units and require the attention of CCU 118. PIU 120 passes these operations on the bus snoop path to CCU 118 indicating a bus snoop and continues to monitor status and response. Other operations are those originated by other units that do not require the attention of the CCU 118. For these operations, PIU 120 only monitors status and response without informing CCU 118.

Clock Distribution and Control 122

[0095] Clock distribution and control 122 contains the logic for gating, shaping, and distributing the internal clocks as well as the off chip cache and directory clocks.

[0096] During normal system operation, all clocks are derived from and synchronized to a single oscillator input by a phase locked loop (PLL) circuit which provides a 'zero delay' clock tree relative to the input oscillator and also a frequency multiplier function. Microprocessor 100 uses this function to run the internal processor logic at a faster rate than the system bus 102 interface logic, which runs at the same rate as the oscillator input. A second on-chip 100 PLL is employed to derive the clocks for the off-chip L2 cache 104. This PLL uses the frequency multiplied output of the first PLL as its input. An off-chip feedback path, constructed to match the path to the cache chips, results in a low skew delay relative to the processor clock domain and allows for synchronous communication between processor 100 and cache 104.

Part 2

Coherency Mechanism for Configurable Caches

[0097] Referring to Figures 11A and 11B, various possible configurations of processors 100 are illustrated interfacing system bus 102. Processors 100...100A, have L2 cache data and tag SRAMs 104...104A attached, and processor 100B does not. In accordance with this invention, a second level cache controller is used in processor 100B in the same manner as if the second level L2 cache existed, such as in the case of L2 cache 104, 104A at processors 100, 100A, and the first level cache status is fed into the second level cache controls, as will be explained hereafter.

[0098] In accordance with the invention, the base microprocessor is not changed to accommodate a design without an L2 cache tag and data arrays. In slightly more detail, second level cache control logic continues to exist within the microprocessor chip. Rather, only the external second level cache tag and data arrays are removed and their inputs to the microprocessor tied to an inactive state. A configuration switch (which feeds controls 470 and 660) is set in second level cache controller 660 that causes bus snoop requests to get reflected onto a first level cache snooping path (the L1 snoop.) The first level cache status is then fed back to the second level cache controller, in a manner consistent with the timing required for support of a second level cache directory search, and fed into the second level cache status signal generation logic, effectively making the second level cache controller believe that the second level cache still exists for the snooping control logic. All other actions remain the same in the second level cache controller and an effective and simple method for supporting snooping bus protocols in the microprocessor has been created. A result is that now every bus request snoops the first level cache without knowledge of presence of an L2 cache. This environment supports entry level single processor configurations where the bus snooping requests only amount to input/output traffic. A similar mechanism is employed for certain requests received from the processor which require cache access. Specifically, these requests can originate from data cache 116 or address translation unit 124. The L1 snooping path is used to determine L1 cache status which is substituted for L2 cache states.

[0099] Referring to Figure 7 in connection with Figure 6, in a system configured as L2 not installed, compare 664, ECC 684 and 686, L2TAGIN register 688 and D/R 694 are not used, and inactivation of these elements is illustrated with hatching. Also, referring to Figure 8 in connection with Figure 5, ECC 632 and 634, L2DATAO register 638, L2DATAI register 640, D/R 644, L2PDO register 636, inpage buffer 646 and L2PDI register 642 are not used in L2 not installed mode, and inactivation of these elements is illustrated with hatching. These blocks represent data and control flow to and from the

L2 tag and data external SRAMs 104, 104A. These external SRAMs are not part of the system in the L2 not installed configuration of processor 100B.

[0100] Referring to Figure 2C and 7, lines 481, 483 to/from data cache control 470 connect into controls 660. These lines are search control 481 and status response 483 and are used to replace L2 status reported from compare 664 when L2 cache is not installed in the system.

[0101] Referring to Figure 7 in connection with Figure 6, DC snoop address generation is fed by lines 679. Lines 679 are generated by multiplexer 678 which is fed by lines 651 and 671. Lines 651 represent a new path used to send a data cache search request address to data cache 116.

[0102] A search operation is one in which the data cache directory is accessed and results are reported for the purpose of determining the status of a cache line in the data cache when the L2 cache is not installed in the system.

[0103] Referring to Figure 10, search operations 575 utilize much of the existing address and control flow which already exists for performing data cache 400 L1 snoop operations with some minor modifications. Searches only report status back to the L2 cache controller 118, replacing the status that would otherwise be generated by accessing the L2 tag array in a system configuration with L2 installed.

[0104] Referring to Figure 10, with L2 installed, pipelined L1 snooping comprises four overlapped stages, as follows: REQUEST 571, SNOOP 572, ACCESS 573 and FLUSH 574. The various registers, arrays and controls comprising these stages have been previously described in connection with Figures 2C, 2D and 6, and are separate pipeline stages from those described with respect to the load/store unit 110, Figure 2A.

[0105] During REQUEST 571, a directory access L1 snoop request is pending to the L1 cache directory 440. If directory address register 414 is available as determined by DC control 470, then the L1 snoop address will be latched into register 414 from cache controller snoop (CCS) address register 670 (Figure 6) on DC snoop address line 679.

[0106] During SNOOP 572, cache directory 440 is accessed and the result of the L1 snoop is latched in DC control 470. At the same time, data cache read address register 408 is latching the address for the access stage of the L1 snoop pipeline from line 415 for access in the following cycle.

[0107] During ACCESS 573, cache arrays 400 are accessed while results from the SNOOP stage are processed. The data read out of cache array 400 during the access stage are latched in register 452.

[0108] During FLUSH 574, the cache data latched in register 452 during the ACCESS stage is sent to L2 CCU over DC write data line 425.

[0109] When data cache 400 is snooped as previously described, any data transfers resulting from the L1

snoop are sent to the L2 CCU 118. L2 CCU 118 in turn transfers this data to either L2 cache 104 or system bus 102, depending on the originator of the L1 snoop request. For instance, a L1 snoop resulting from an ATU 124 request will cause the data being transferred by the snoop operation to be placed into L2 cache 104. If, on the other hand, the L1 snoop is resulting from a system bus operation the data will be transferred out to system bus 102. L1 snoop data from data cache L1 400 may be merged with L2 cache 104 data so that an entire 128 bytes corresponding to an L2 cache 104 line size will be transferred out to the system bus 102. When placed on system bus 102, the data will then either be directed to main memory 126 or another processor 100B, depending on the operation.

[0110] If multiple data transfers are required out of L1 cache 400 for a given cache line, then a pipeline hold is sent to the REQUEST 571 and SNOOP 572 stages of the pipeline and the ACCESS 573 stage is then repeated.

[0111] Referring further to Figure 10, with L2 cache not installed, the mechanism for doing an explicit search request in data cache 400 is similar to that for a L1 snoop. The main differences are that no data transfers ever directly result from a search and that only one search request is done at a time. The REQUEST 571 and SNOOP 572 stages of a pipelined L1 snoop operation are essentially reused for performing a search 575. The results of a search 575 are reported to L2 cache controller 118 in the cycle that would correspond to an access 573. That is, in the case of a search 575 the data cache 400 is not accessed but the results of the snoop stage 572 are still processed and then sent to L2 cache controller 118.

[0112] Referring to Figures 12 and 13, there are two types of search requests:

1) Explicit searches, where the status in the data cache directory 440 is interrogated by sending a DC snoop address from L2 cache controller 118 to data cache 400.

2) Implied searches, where the status of the data cache directory 440 is not directly interrogated, but the search status response lines 483 are checked by L2 cache controller 118. The data cache controller 470 is responsible for ensuring that these lines 483 indicate the appropriate state for implied searches.

[0113] As an example, an ATU 124 request to fetch a page table for address translation results in an explicit search 575 of the data cache 400 in L2 not installed configurations. The page table may be resident in a line stored in data cache 400. If a search results in a data cache miss, the operation proceeds out to main memory 126 to fetch the page tables. If the search results in a data cache hit, a separate L1 snoop operation 579 oc-

curs to data cache 400 to flush out any modified data. L1 snoop operation 579 includes request 571, snoop 572, access 573 and flush 574. The L1 snoop operation 579 resulting from the search 575 is the same mechanism that is used in a system with L2 cache. If the page table is resident in data cache 400, a system with L2 cache would have determined an L2 hit from accessing the L2 tag array. A L1 snoop operation would then be sent to the data cache.

[0114] For either an implied search or an explicit search, when the L2 cache controller 118 takes a DC or ATU request, it initiates operation of a tag state machine 708 (Figure 9) as if an L2 tag access is occurring. When the tag state machine reaches a point where L2 status would otherwise be available, data cache search status is queried instead.

[0115] An example of an implied search is a DCU 116 request for data due to a data cache 400 miss. In this case, the data cache 400 status is already known and another directory access via a search operation 575 is not necessary. Data cache 400 indicates miss status on the status response line 483 in the appropriate cycle. L2 cache controller 118 then initiates an operation to access main memory 126 for the requested data in a manner similar to a system with an L2 cache that has an L2 cache miss.

[0116] Another case of an implied search is a DCU 116 request for exclusive access to a cache line, otherwise known as a Dclaim request. This request is initiated when data cache 400 currently has a cache line in the shared state and has a store pending to that line. The L2 cache controller 118 starts the tag state machine and checks data cache search status as in the previous case. In this case, however, the data cache 400 reports shared status on the search status lines 483. The L2 cache controller 118 then initiates a Dclaim on the system bus in a manner similar to a system with L2 cache that has the line in the shared state.

[0117] Another case of an explicit search is a bus snoop request. The bus snoop request originates from a system bus command latched up in processor interface unit (PIU) 120. This bus snoop is then forwarded on to L2 cache controller 118. The L2 cache controller 118 starts its tag state machine and sends a search request, and as previously described, L1 status lines 483 is substituted for L2 status lines 665. This status 483 is then forwarded to the system bus 102 interface from control block 660 to PIU 120 and thence to system bus 102. Based on the status and the type of bus snoop, a separate L1 snoop operation 579 may then be initiated. Data may need to be flushed from the L1 cache, in which case it is transferred to the L2 cache data flow along the path 425 to CCDI 624 and then on store data lines 625 to data registers 502 and 504. These data are then transferred to system bus 102 on lines 513 and 517 through D/R blocks 512 and 516.

[0118] Referring to Figures 12 and 13, in all of the search examples given above, an access of the L2 tag

array is replaced with either an explicit or implicit search operation. If data needs to be transferred out of data cache 400 due to the results of a search, then a separate L1 snoop operation 579 is sent to the data cache 400.

[0119] Referring to Figure 9 in connection with Figure 6 and the timing charts of Figures 12 and 13, a state diagram for the tag state machine 708, which is a portion of controls 660, is set forth. As is represented by arrow 703, tag state machine 708 watches in tag access cycle 1 tacc1[0] state 702 for a command from address/command register 650 on lines 651. In systems where L2 cache is installed, controls 660 tag access cycle 1, state tacc1[0] 702, results in an address being sent on lines 693 to the L2 tag array (also referred to as the L2 cache directory) within cache 104 based on the command received on lines 651. Controls 660 tag access cycle 2, state tacc2[1] 704, represents the timing cycle for off chip pipelined SRAMS that implement the tag arrays; tag data is latched in L2 cache controller 118 in L2 tag in register 688 at the end of this state 704. In controls 660 tag compare cycle, state tcmpr[2] 706, L2 tag data is available for determining an L2 cache hit.

[0120] In a system with L2 cache not installed, tacc1[0] 702 responds by activating search control line 481 and also a search address is sent from lines 651 through multiplexer 678 to lines 679 and then to L1 directory address register 414. The L1 cache directory is then accessed in the cycle where tacc2[1] 704 is active. L1 cache directory 440 status 483 is substituted in this state tcmpr[2] 706 for L2 directory hit status.

[0121] A system with an L2 cache in the preferred embodiment implements an 128 byte L2 cache 104 line size and a 64 byte data cache 400 line size. All memory coherence protocols operate on the L2 cache line size of 128 bytes. In a system with L2 cache not installed, a configuration bit is changed such that all memory coherence protocols operate on the L1 cache line size of 64 bytes. In accordance with another embodiment, in order to maintain cache coherency in a mixed system having processors 100, 100A with L2 104, 104A installed and processors 100B without L2 installed, all on the same system bus 102, the L1 and L2 cache line sizes are kept the same.

[0122] Data cache directory 440 has a single read port in the preferred embodiment and as a result, search operations 575 must present a busy condition to the processor so that cache directory 440 access can occur. In certain circumstances, a search request encounters a processor operation in progress that requires access to data cache directory 440. In this circumstance a busy is presented to the L2 cache controller 118 on status lines 483 to indicate that the search request cannot be honoured immediately. Once the busy drops to the L2 cache controller, the search operation is initiated.

[0123] An error can also be detected in the data cache directory 440 while a search operation 575 is in progress. In this case, search response status 483 indicates that the search operation should be retried. If the

search request results from a system bus operation, then the L2 cache controller will issue a response on the system bus that the operation should be retried.

5 Advantages over the Prior Art

[0124] It is an advantage of the invention according to the preferred embodiment that a method and system is provided for achieving data cache consistency in a multiprocessor or uniprocessor system including either entry level processors having no L2 cache, or high level processors including an L2 cache.

[0125] It is a further advantage of the system that such is achieved with a common microprocessor chip design.

[0126] It is still a further advantage of the system that a processor design including provision for control of both an L1 cache and an installed L2 cache properly functions when the L2 cache is not installed.

20 Alternative Embodiments

[0127] It will be appreciated that, although specific embodiments of the invention have been described herein for purposes of illustration, various modifications may be made without departing from the spirit and scope of the invention.

Claims

1. A cache controller for the management of contents of a single level of cache, comprising:

a second level cache controller including a first level cache snooping path; and

a first system interface; wherein

said second level cache controller is operable in the absence of a second level cache for reflecting bus snoop requests received at said interface to said first level cache snooping path.

2. The cache controller of claim 1, wherein

the second level cache controller includes status generation logic for generating a first status signal; the cache controller further comprising:

a first level cache controller for generating and feeding to said status generation logic a second status signal consistent with timing required for support of a second level cache directory access;

whereby said second level cache controller is enabled to operate and respond as though an access of a second level cache directory had

occurred.

3. The cache controller of claim 2, wherein said first level cache controller is responsive to a snoop request forwarded from said second level cache controller for conducting a search operation to generate said second status signal.
4. The cache controller of claim 3, wherein said first system interface is an interface to a system bus and responses to bus snoop requests are directed to said system bus for delivery selectively to another processor or to another device on said system bus.
5. The cache controller of claim 3, said search operation comprising an explicit search.
6. The cache controller of claim 5, further comprising:
 - a data cache directory;
 - said second level cache controller being further operable for executing said explicit search by interrogating said data cache directory with a data cache snoop address and then checking said second status signal.
7. The cache controller of claim 6, wherein first system interface is a system bus interface, further comprising:
 - a processor interface unit selectively operable for generating a bus snoop request;
 - a tag state machine;
 - said second level cache controller being further operable responsive to said bus snoop request for:
 - executing said explicit search to determine if data associated with said bus snoop request resides in a cache line stored in said first level cache and selectively generate said second status signal as a cache miss or cache hit signal; and
 - initiating operation of said tag state machine for querying said second status signal, and:
 - responsive to generation of a first level cache miss signal, returning status to said system bus interface; and
 - responsive to generation of a first level cache hit signal, causing a separate L1 snoop operation to said first level cache to flush out any modified data, and forwarding said modified data

ta to said system bus.

8. A cache controller according to any preceding claim, wherein said first system interface is a system bus interface, and wherein the cache controller includes a processor interface; and wherein said second level cache controller is operable in the absence of a second level cache for reflecting to said first level cache snoop path processor snoop requests received at said processor interface and bus snoop requests received at said system bus interface.
9. The cache controller of claim 3, wherein said first level cache controller is further operable to present a busy condition to said processor during search operations.
10. The cache controller of claim 9, said second level cache controller being responsive to a busy signal from said first level cache controller indicating a processor operation in progress that requires access to a data cache directory for holding said search operation pending until said busy signal terminates and thereafter for initiating said search operation.
11. The cache controller of claim 3, said second level cache controller being further operable responsive to an error being detected during said search operation for causing said search operation to be retried.
12. The cache controller of claim 3, said search operation selectively comprising an explicit search or an implicit search.
13. The cache controller of claim 12, further comprising:
 - a data cache directory;
 - said second level cache controller being further operable for executing said explicit search by interrogating said data cache directory with a data cache snoop address and then checking said second status signal; and
 - said second level cache controller being further operable for executing said implicit search by checking said second status signal.
14. The cache controller of claim 13, further comprising:
 - an address translation unit selectively operable for generating an address translation page table fetch request;
 - a tag state machine;

said second level cache controller being further operable responsive to said fetch request for:

executing said explicit search to determine if said page table resides in a cache line stored in said first level cache and selectively generate a said second status signal as a cache miss or a cache hit signal; and

initiating operation of said tag state machine for querying said second status signal, and:

responsive to generation of a first level cache miss signal, fetching said page table from main memory; and

responsive to generation of a first level cache hit signal, causing a separate L1 snoop operation to said first level cache to flush out any modified data.

15. The cache controller of claim 13, further comprising:

a data cache unit selectively operable for issuing a data request due to a data cache miss;

a tag state machine;

said second level cache controller being further responsive to said data request for initiating operation of said tag state machine for querying said second status signal and, responsive to said second status signal indicating a data cache miss, for initiating a main memory access for the requested data.

16. The cache controller of claim 13, further comprising:

a data cache unit selectively operable responsive to said first level cache having a cache line in a shared state and a store pending to that cache line for issuing a Dclaim request for exclusive access to said cache line;

a tag state machine;

said second level cache controller being operable for initiating operation of said tag state machine querying said second status signal and, responsive to said second status indicating shared status, for initiating a DClaim.

17. A cache controller according to any one of claims 1 to 3 wherein the first system interface is an interface to a processor; and wherein

said second level cache controller is operable in the absence of a second level cache for reflecting processor snoop requests received at said interface

to said first level cache snooping path.

18. The cache controller of claim 3, wherein responses to processor snoop requests result in commands that are selectively directed to a system bus for delivery to main memory or another processor, or to said first level cache controller for performing a separate snoop operation.

19. A method for managing contents of a single level cache, comprising the steps of:

reflecting second level cache snoop requests received at a second level cache controller to a first level cache controller cache snooping path;

responsive to a reflected snoop request, operating said first level cache controller to generate and feed to said second level cache controller a status signal consistent with timing required for support of a second level cache directory access.

20. Cache coherency apparatus for a system including one or more processor nodes, a system bus interconnecting said processor nodes and other system devices, with all processor nodes operating without an L2 cache installed, the cache coherency apparatus at said one processor node comprising:

a second level cache controller including a first level cache snooping path;

a processor interface to said second level cache controller;

said second level cache controller including status generation logic for responding with a first status signal to a processor snoop request message on said processor interface directed to a second level cache;

means for reflecting said processor snoop request to said first level cache snooping path;

a first level cache controller operable responsive to said reflected processor snoop request for generating and feeding to said status generation logic a second status signal consistent with timing required for support of a second level cache directory access; and

said second level cache controller further being responsive to said second status signal for replacing said first status signal.

21. Cache coherency apparatus for a system including

one or more processor nodes, a system bus interconnecting said processor nodes and other system devices, with all processor nodes operating without an L2 cache installed, the cache coherency apparatus at said one processor node comprising:

5

a second level cache controller including a first level cache snooping path;

said second level cache controller including status generation logic for responding with a first status signal to a bus snoop request message on said system bus directed to a second level cache;

10
15

means for reflecting said bus snoop request to said first level cache snooping path;

a first level cache controller operable responsive to said reflected bus snoop request for generating and feeding to said status generation logic a second status signal consistent with timing required for support of a second level cache directory access; and

20
25

said second level cache controller further being responsive to said second status signal for replacing said first status signal.

22. Cache coherency apparatus for a system including a plurality of processors, a system bus interconnecting said processors, and at least one processor node without an L2 cache installed, the cache coherency apparatus at said processor node comprising:

30
35

a second level cache controller including a first level cache snooping path;

a processor interface to said second level cache controller;

40

said second level cache controller including status generation logic for responding with a first status signal to a bus snoop request on said system bus directed to a second level cache or to a processor snoop request on said processor interface;

45

means for reflecting said bus snoop request or said processor snoop request to said first level cache snooping path;

50

a first level cache controller operable responsive to said reflected bus snoop request or processor snoop request for generating and feeding to said status generation logic a second status signal consistent with timing required for

55

support of a second level cache directory access; and

said second level cache controller further being responsive to said second status signal for replacing said first status signal.




FIG. 1




FIG. 2




FIG. 2A




FIG. 2B




FIG. 2C




FIG. 2D




FIG. 3




FIG. 4




FIG. 5




FIG. 6




FIG. 7




FIG. 8




FIG. 9




FIG. 10




FIG. 11A




FIG. 11B

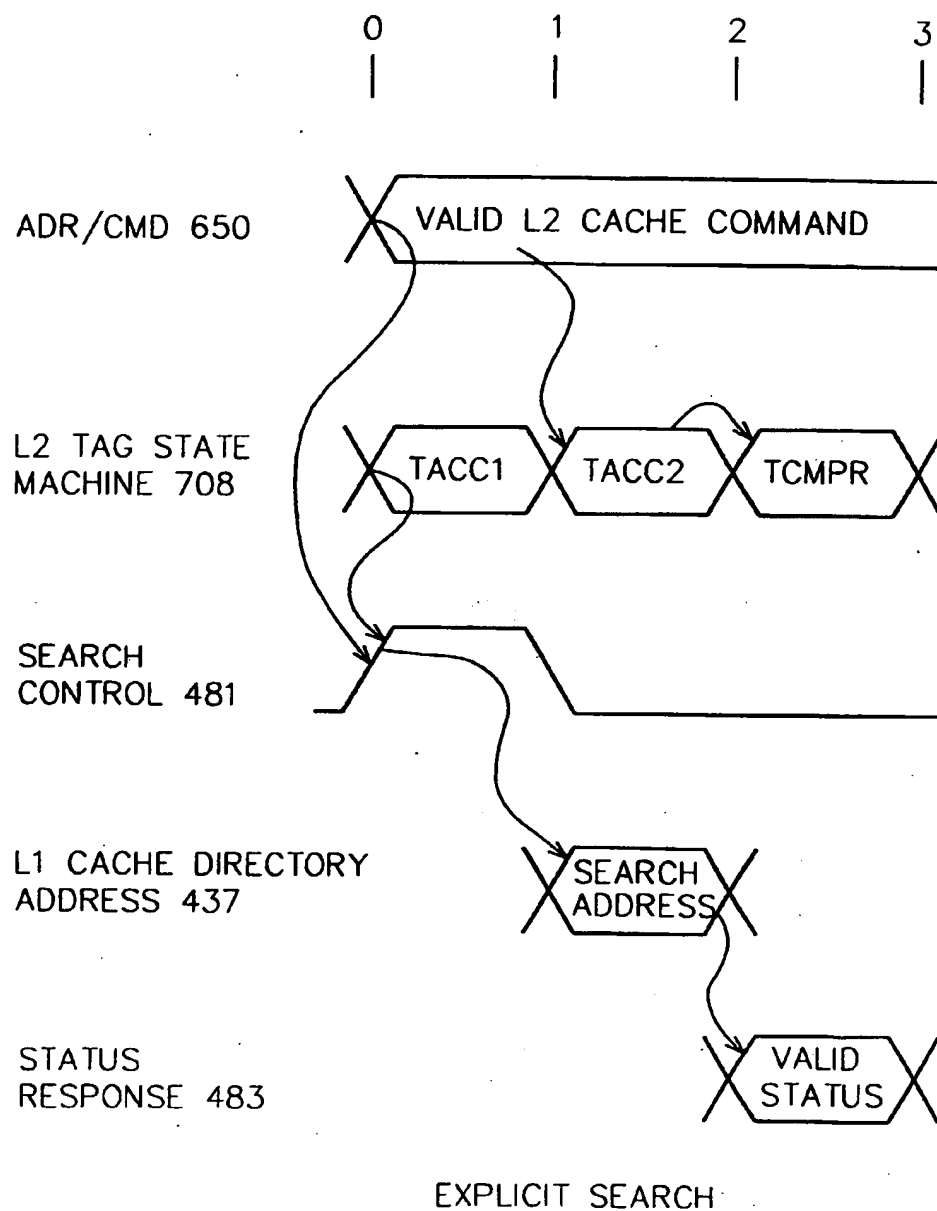



FIG. 12



IMPLIED SEARCH

FIG. 13



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 98 30 5910

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.6)
X	GB 2 200 481 A (NAT SEMICONDUCTOR CORP) 3 August 1988	1,17	G06F12/08
A	* page 33, line 28 - page 34, line 12; figures 21,22 *	19-22	
A	EP 0 743 601 A (IBM) 20 November 1996 * column 4, line 49 - column 5, line 28 *	1-22	
A	LEVITAN D ET AL: "THE POWERPC 620TM MICROPROCESSOR: A HIGH PERFORMANCE SUPERSCALAR RISC MICROPROCESSOR" DIGEST OF PAPERS OF THE COMPUTER SOCIETY COMPUTER CONFERENCE (SPRING) COMPCON, TECHNOLOGIES FOR THE INFORMATION SUPERHIGHWAY SAN FRANCISCO, MAR. 5 - 9, 1995, no. CONF. 40, 5 March 1995, pages 285-291, XP000545441 INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS * page 291, left-hand column, paragraph 4.3 *	1-22	
A	EP 0 533 373 A (NCR INT INC) 24 March 1993 * column 4, line 1 - column 5, line 18; figures 2-4 *	1,19-22	TECHNICAL FIELDS SEARCHED (Int.Cl.6) G06F
The present search report has been drawn up for all claims			
Place of search THE HAGUE		Date of completion of the search 13 November 1998	Examiner Nielsen, O
CATEGORY OF CITED DOCUMENTS X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document			

EPO FORM 1503 (3.92) (P4/D01)

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- ☐ BLACK BORDERS
- ☐ IMAGE CUT OFF AT TOP, BOTTOM OR SIDES
- ☐ FADED TEXT OR DRAWING
- ☐ BLURRED OR ILLEGIBLE TEXT OR DRAWING
- ☐ SKEWED/SLANTED IMAGES
- ☐ COLOR OR BLACK AND WHITE PHOTOGRAPHS
- ☐ GRAY SCALE DOCUMENTS
- ☐ LINES OR MARKS ON ORIGINAL DOCUMENT
- ☒ REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY
- ☐ OTHER: _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.